



Scalable Coordination Architectures For Deeply Distributed Systems (SCADDS)

Final R & D Report

Period: 29 June 1999 – 31 March 2003

*Defense Advanced Research Projects Agency (DoD)
Information Technology Office (ITO)
Under Grant # DABT63-99-1-0011
Issued by Directorate of Contracting
Fort Huachuca, AZ*

BEST AVAILABLE COPY

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

20040721 092

REPORT DOCUMENTATION PAGE				FORM APPROVED OMB NO. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503					
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 08-07-2004		2. REPORT TYPE Final Technical Report		3. DATES COVERED (From - To) 29 June 1999 to 31 March 2003	
4. TITLE AND SUBTITLE Scalable Coordination Architecture for Deeply Distributed Systems (SCADDS)				5a. CONTRACT NUMBERS	
				5b. GRANT NUMBER DABT63-99-1-0011	
				5c. PROGRAM ELEMENT NUMBER	
				5d. PROJECT NUMBER	
6. AUTHOR(S) Deborah Estrin, John Heidemann, and Ramesh Govindan				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) USC INFORMATION SCIENCES INSTITUTE 4676 ADMIRALTY WAY MARINA DEL REY, CA 90292-6695				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Directorate of Contracting ATTN: ATZS-DKO-I Post Office Box 12748 Fort Huachuca, AZ 85670-2748				10. SPONSORING/MONITOR'S ACRONYM(S)	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER	
12. DISTRIBUTION AVAILABILITY STATEMENT UNCLASSIFIED/UNLIMITED					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The SCADDS project explored and demonstrated scalable coordination mechanisms for deeply distributed and dynamic systems. Nodes in these systems were heterogeneous, having a range of sensing, actuation and communication capabilities. These systems raised many challenges for distributed system and network design. The first was a shift from node-centric, to data-centric network architecture. Both scalability and long lifetime called for extensive processing of data within and among the nodes of the sensor network. Rather than streaming all sensor readings back to a central site for processing, nodes autonomously exchanged data, filtered out uninteresting events, and identified patterns of interest. The second challenge was to build systems that were truly self-configuring; able to adapt efficiently to ad hoc deployment and both environmental and network dynamics. This paradigm shift required new network architecture.					
15. SUBJECT TERMS SCADDS, GEAR, S-MAC, Diffusion, Sensor Network					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UNCLASSIFIED	18. NUMBER OF PAGES 23	19a. NAME OF RESPONSIBLE PERSON John Heidemann
a. REPORT UNCLASSIFIED	b. ABSTRACT UNCLASSIFIED	c. THIS PAGE UNCLASSIFIED			19b. TELEPHONE NUMBER (310) 822-1511

Scalable Coordination Architecture for Deeply Distributed Systems (SCADDS)

Table of Contents

1. INTRODUCTION.....	1
2. DIRECTED DIFFUSION.....	2
2.1 THE PUBLISH/SUBSCRIBE API AND DATA NAMING.....	3
2.2 DIRECTED DIFFUSION PROTOCOL FAMILY	4
2.3 COMPARING DIFFUSION WITH ALTERNATIVES	5
2.4 TINY DIFFUSION.....	6
2.5 GEOGRAPHICAL AND ENERGY AWARE ROUTING (GEAR).....	7
3. SENSOR-MAC PERFORMANCE IN MULTI-HOP NETWORKS.....	8
4. APPLICATION TECHNIQUES.....	4
4.1 CLUSTERING	10
4.2 AGGREGATION	10
4.3 TOPOLOGY CONTROL.....	11
5. INFRASTRUCTURE	5
5.1 TESTBED DEVELOPMENT	13
5.2 TESTBED – IMPLEMENTATION AND ANALYSIS	13
5.2.1 <i>Connectivity Measurements for SCADDS Testbed</i>	13
5.2.2 <i>Experiment Details</i>	14
5.2.3 <i>Results</i>	14
5.3 SOFTWARE DEVELOPMENTS	15
5.4 SENSIT COLLABORATION/INTEGRATION	16
6. PERSONNEL.....	6
7. REFERENCES	7
7.1 WHITE PAPERS/TECHNICAL REPORTS	17
7.2 CONFERENCE/JOURNAL	18

Table of Figures

Figure 1: Modules using filter architecture with several variants of directed diffusion	6
Figure 2: Average packet latency on each hop in the lowest traffic load	9
Figure 3: A Virtual Grid	11
Figure 4: Compares GAF and CEC performance.....	12
Figure 5: Shows experimental performance of CEC in UCLA's testbed.....	12
Figure 6: Hardware platforms supporting diffusion.....	13
Figure 7: Experimentally measured connectivity in testbed	15

Scalable Coordination Architectures for Deeply Distributed Systems Final Report 29 June 1999 – 31 March 2003

1. Introduction

This report summarizes the progress of the entire SCADDS project at ISI from June 29, 1999 to March 31, 2003. Over this period, the SCADDS project was overseen by the principal investigators: Deborah Estrin, John Heidemann, and Ramesh Govindan.

The SCADDS project explored and demonstrated scalable coordination mechanisms for deeply distributed and dynamic systems. Nodes in these systems were heterogeneous, having a range of sensing, actuation and communication capabilities. These systems raised many challenges for distributed system and network design. The first was a shift from node-centric to data-centric network architecture. Both scalability and long lifetime called for extensive processing of data within and among the nodes of the sensor network. Rather than streaming all sensor readings back to a central site for processing, nodes autonomously exchanged data, filtered out uninteresting events, and identified patterns of interest. The second challenge was to build systems that were truly self-configuring; able to adapt efficiently to ad hoc deployment and both environmental and network dynamics. This paradigm shift required new network architecture. The SCADDS project investigated an approach called directed diffusion. The report is divided into the following subject areas:

Directed Diffusion Architecture

Diffusion Architecture:

Diffusion is a tunable algorithm for information dissemination in wireless sensor networks. Rapid deployment of large numbers of sensors in dynamic and potentially hostile environments presents a challenge to existing networking techniques, in terms of scalability, robustness, and adaptability. The SCADDS project explored the use of localized algorithms as a building block for such dynamic data-dissemination systems.

Tiny Diffusion

Full diffusion is designed for 32-bit computer. Tiny Diffusion is a subset of Full Diffusion that runs on Berkley motes.

Geographical and Energy Aware Routing (GEAR)

Geographic and Energy Aware Routing (GEAR) is an extension to diffusion that uses energy aware and geographically informed neighbor selection to route a packet towards the target region. This strategy attempts to balance node energy consumption and thereby increase network lifetime. Within a region, it uses a recursive geographic forwarding technique or a restricted flooding algorithm to disseminate the packet.

Sensor-MAC (S-MAC)

The S-MAC is an energy-efficient MAC protocol explicitly designed for wireless sensor networks. Its major goal is to reduce energy consumption while maintaining good properties such as collision avoidance, scalability, and self-organization. The protocol reduces energy consumption through the following approaches: reduce idle listening, avoid collisions, avoid overhearing and reduce control overhead.

Application Techniques

Complementing network-level advances, the SCADDS project explored several application-level techniques to improve sensor net performance.

Clustering

Clustering can contribute to more scalable system behavior as number of sensors increase, improved robustness, and efficient resource utilization for many distributed sensor coordination tasks. However, self-configuring techniques required to organize sensors into clusters (since manual configuration of a large number of sensors is not feasible) can consume significant resources that need to be amortized over the gains in application function. The SCADDS project investigated this trade-off with energy as the primary resource constraint in the context of an object tracking sensor network application. In this application, sensors monitor location and status of various tagged objects (e.g., projectors, cameras). Queries from users about these tagged objects were efficiently resolved through a hierarchical scheme.

Aggregation

Aggregation can be expressed as local transformation of data, to reduce overall power consumption. Hence, it might take various forms: from compressing data, to merging data temporally or spatially (i.e. data coming from different neighbors), to only forwarding deductions on the basis of received data.

Topology Control

Topology control protocols extend network lifetime by periodically turning node or node radios off. The sleeping time of radio interface can be adapted according to the nodes' density. The idea uses network density as a clue to put some nodes into sleep mode. The design takes advantage of the existing ad hoc routing protocols to find network neighborhood information, and adapts to the node's sleep time according to the number of its neighbors. The neighborhood information is stored in soft state so that the algorithm can adapt to the network dynamics. The document explains several different topological control protocols BECA, AFECA, GAF, and CEC, both in simulation and experimentation.

Infrastructure and Collaboration

An emphasis as part of SensIT was on providing research and tools that can be used by other researchers.

SCADDS was very active at integration and SensIT field tests at 29 Palms.

Testbed Development

The ISI testbed provides an experimentation environment to study communication protocols in a large network with different topologies. With feedback from the well-instrumented real-world experiments, one can iteratively improve performance of protocols and algorithms and validate simulation results. Moreover, playing with real systems can lead us to better understanding of the design challenges.

Software

The project developed many software tools over the course of this project, both as research itself and as tools to support the research. These include: directed diffusion with support for hardware Sensoria, WINSng 1 and 2, Ethernet evaluation, Radiometrix, mote radios, extensions in aggregation GEAR, PUSH, and nested queries. Diffusion was also posted to run inside the ns-2 simulator. The project also developed these utilities: Emlog and Parapin Diffusion visualization.

2. Directed Diffusion

Directed diffusion [Intanagonwiwat00a] is a data-centric data dissemination protocol. Data generated by sensor nodes is named by attribute-value pairs. A node requests data by sending interests for named data. Data matching the interest is then "drawn" down towards that node. Intermediate nodes can cache, or transform data, and may direct interests based on previously cached data (Section 3).

Directed diffusion is significantly different from IP-style communication where nodes are identified by their end-points, and inter-node communication is layered on an end-to-end delivery service provided within the network. In directed diffusion, nodes in the network are application-aware as we allow application-specific code to run in the network and assist diffusion in processing messages. This allows directed diffusion to cache and process data in the network (aggregation), decreasing the amount of end-to-end traffic, and resulting in higher energy savings. We show that using directed diffusion one can realize robust multi-path delivery, empirically adapt to a small subset of network paths, and achieve significant energy savings when intermediate nodes aggregate responses to queries (Section 5).

2.1 The Publish/Subscribe API and Data Naming

Directed diffusion uses a publish/subscribe-based API. To receive data, users or programs subscribe to a particular set of attributes, becoming data sinks. A callback function is then invoked whenever relevant data arrives at the node. Sensors publish data that they have, becoming data sources. In both cases, what data is provided or received is described by an attribute-based naming scheme described next. It is the job of the diffusion dissemination algorithms to ensure that data is communicated efficiently from sources to sinks across a multi-hop network. In general, publishing and subscribing sends messages across the network. The exact cost of these operations depends on which diffusion algorithm is used.

To allow applications to influence data as it moves through the network, users can create filters at each sensor node with the filter APIs. Filters indicate what messages they are interested in by attributes; each time a matching message arrives at that node the filter is allowed to inspect and alter its progress in any way. Filters can suppress messages, change where they are sent next, or even send other messages in response to one (perhaps triggering further sensors to satisfy a query).

Diffusion uses an attribute-based naming scheme to associate sources and sinks and to trigger filters. This flexible approach to naming is important in several ways. First, attribute-based naming is consistent with the publish/subscribe application-level interface (just described) and many-to-many communication. Diffusion's naming scheme is data-centric, allowing applications to focus on what data is desired rather than on individual sensor nodes. The approach also supports multiple sources and sinks, rather than simple point-to-point communication. Thus applications may subscribe to "seismic sensors in the southeast region" rather than seismic sensors #15 and #35, or hosts 10.1.2.40 and 10.2.1.88.

Second, diffusion attributes provide some structure to a message. By identifying separate fields, data dissemination algorithms can use application data to influence routing. For example, application-specific, geographic information can limit where diffusion must look for sensors. In addition, treating messages as sets of attributes simplifies application and protocol extensions (a need also suggested for future Internet-based protocols).

Finally, attributes serve to associate messages with sources, sinks, and filters via matching. If the attributes in a sink's subscription match those of source's publication, diffusion must send any published data to the sink.

2.2 Directed Diffusion Protocol Family

Publish/subscribe provides an application's view to a sensor network, and attribute-based naming a detailed way to specify which sources and sinks communicate. The "glue" that binds the two are the directed diffusion algorithms for data dissemination. In a traditional network, communication is effected by routing, usually based on global addresses and routing metrics. Instead, we use the term data dissemination to emphasize the lack of global addresses, reliance on local rules, and the use of application-specific in-network processing.

The original, two-phase directed diffusion uses several control messages to realize our publish/subscribe API: sinks send interest messages to find sources, sources use exploratory data messages to find sources, and positive and negative reinforcement messages select or prune parts of the path. Early work [Intanagonwiwat00a] identified these primitives, described the concept of diffusion, and evaluated a specific algorithm that we now call two-phase pull diffusion. We found this algorithm ideal for some applications but as our experience with sensor networks applications grew, we found two-phase pull a poor match for other classes of applications.

We see diffusion not as a single algorithm, but as a family of algorithms built from these primitives. Other algorithms provide better performance for some

applications. We have recently made two additions to the diffusion protocol family: one-phase push and one-phase pull [Heidemmann03b].

Over the course of the project, SCADDS researchers developed several verresionons of diffusion with different performance characteristics and evaluated them through analysis, simulation, and experimentation. See [Intanagonwiwat03a] for the primary analytic results, [Intanagonwiwat00a] and [Heidemmann03b] for simulation results evaluating the basic algorithm and variants, and [Heidemmann01c] for experimental results. We summarize the key results below.

2.3 Comparing diffusion with alternatives

Our first experiment compares diffusion to omniscient multicast and flooding scheme for data dissemination in networks. Figure 6(a) shows the average dissipated energy per packet as a function of network size. Omniscient multicast dissipates a little less than a half as much energy per packet per node than flooding. It achieves such energy efficiency by delivering events along a single path from each source to every sink. Directed diffusion has noticeably better energy efficiency than omniscient multicast. For some sensor fields, its dissipated energy is only 60% that of omniscient multicast. As with omniscient multicast, it also achieves significant energy savings by reducing the number of paths over which redundant data is delivered. In addition, diffusion benefits significantly from in-network aggregation. In our experiments, the sources deliver identical location estimates, and intermediate nodes suppress duplicate location estimates. This corresponds to the situation where there is, for example, a single vehicle in the specified region.

Figure (1) plots the average delay observed as a function of network size. Directed diffusion has a delay comparable to omniscient multicast. This is encouraging. To a first approximation, in an uncongested sensor network and in the absence of obstructions, the shortest path is also the lowest delay path. Thus, our reinforcement rules seem to be finding the low delay paths. However, the delay experienced by flooding is almost an order of magnitude higher than other schemes. This is an artifact of the MAC layer: to avoid broadcast collisions, a randomly chosen delay is imposed on all MAC broadcasts. Flooding uses MAC broadcasts exclusively. Diffusion only uses such broadcasts to propagate the initial interests. On a sensor radio that employs a TDMA MAC-layer, we might expect flooding to exhibit a delay comparable to the other schemes.

In summary, directed diffusion exhibits better energy dissipation than omniscient multicast and has good latency properties.

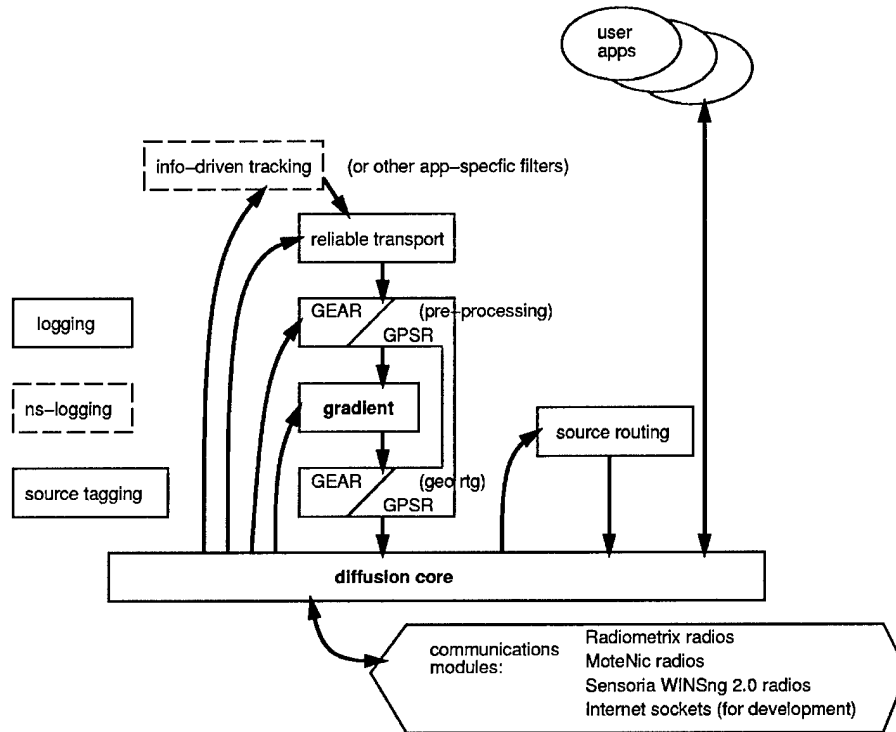


Figure 1: Modules using filter architecture with several variants of directed diffusion

2.4 Tiny Diffusion

TinyDiffusion is implemented as a component in TinyOS that adds 3250B code and 144B of data (including support for radio and a photo sensor), so the entire system runs in less than 5.5KB of memory.

TinyDiffusion is statically configured to support for 5 active gradients and a cache of 10 packets of the 2 relevant bytes per packet. Although reduced in size, the logical header format is compatible with that of the full diffusion implementation and SCADDS is implementing software to gateway between the implementations. Although SCADDS does not currently provide filters in tiny-diffusion, they are an essential component of enabling in-network aggregation in diffusion, and SCADDS intends to add them. The team intends to leverage on the ability to reprogram motes over the air to program filters dynamically.

Motes and tiny-diffusion can be used in regions where there is need for dense sensor distribution, such as distributing photo sensors in a room to detect change in light or temperature sensors for fine grained sensing. They provide the necessary sensor data processing capability, with the ability to use diffusion to communicate with less resource constrained nodes (PC/104-class nodes with the full system). They can also be used to provide additional multi-hop capability under adverse wireless communication conditions.

We thus envisage deployment of a tiered architecture in which less resource constrained nodes running the full system will form the highest tier, and act as gateways to the second tier comprising motes connected to low-power sensors running tiny-diffusion. Most of the network “intelligence” is programmed into the

first tier. Aggregation filters that are determined by the full nodes are programmed onto the motes over the air.

2.5 Geographical and Energy Aware Routing (GEAR)

SCADDS designed and studied (through simulation) the proposed Geographic and Energy Aware Routing (GEAR) algorithm, which uses energy aware and geographically informed neighbor selection to route a packet towards the target region. This strategy attempts to balance node energy consumption and thereby increase network lifetime. Within a region, it uses a recursive geographic forwarding technique or a restricted flooding algorithm to disseminate the packet.

SCADDS simulated the above algorithm for uniform and non-uniform traffic distributions, and compared its performance to a non energy-aware geographic routing algorithm, GPSR. For non-uniform traffic, GEAR delivers 70% to 80% more packets than GPSR. For uniform traffic, GEAR successfully delivers 25 - 35% times more packets than GPSR. Hence, GEAR exhibits more gain in non-uniform traffic scenarios than uniform traffic scenarios. The explanation is that when traffic sources are clustered together, GEAR's energy balancing efforts pay off most. However, in both cases, the GEAR algorithm performs better in terms of connectivity after partition (i.e., how routing the given traffic patterns affects the rest of the network) and the number of normalized traffic pairs broken down per delivered packet.

SCADDS studied how the protocol's performance being sensitive to imprecise neighbor information. Simulation results show that it is not necessary to update neighbor information for every packet. With increasing update threshold, the protocol performance degrades gracefully, but the number of control packets generated drops dramatically.

SCADDS currently implement a prototype of GEAR protocol in a moderate size testbed, and are going to test it in real world environment. The team is also in the process of porting GEAR to the diffusion implementation in ns-2. On the other hand, not all applications always require the same service in terms of delay, quality of data, or energy cost. For instance, most of the time, sensor net publications would trade delay for energy efficiency. However, sometimes, it may be concerned more about fast response than energy cost or the quality of data. There is no previous QoS work in sensor network context. Although there have been lots of QoS research for the Internet, the unique characteristics and constraints in sensor net pose new challenges, and make the problem somewhat different from its counterpart in the Internet domain.

First of all, because of its stringent energy constraint a sensor net tends to be more energy constrained than bandwidth constrained. This makes admission control, scheduling, queue management not the focus of ToS in sensor net any more, although it may still remain an issue in some sensor net application context. Second, Because of its scarce energy resources, sensor networks cannot afford over-provisioning or brute-force approaches. Thirdly, deployment in large numbers requires the algorithm to be fully distributed, the question is how to achieve end-to-end flow characteristics with a distributed local algorithm.

Before jumping into any solution, SCADDS used a simple analytical model to show that the gain of ToS is maximized when sensor network is densely deployed,

thus enabling multiple alternative paths; when the delay sensitive traffic and regular traffic is mixed, especially when delay sensitive traffic is only a small percentage of total traffic, so that all of them can be accommodated on the paths that can satisfy the delay requirements; and when delay-sensitive applications have strict delay requirements. The team is going to study the design space of ToS in sensor networks, and design algorithms that can satisfy application requirements and conserve network energy at the same time.

Previously, SCADDS tested Geographical and Energy Aware Routing (GEAR) in a high level simulator, which ignored many low level details. In order to test the protocol with more realistic lower layer details, the team implemented and evaluated the protocol in ns-2. In the ns simulation, the team validated the two major design choices of GEAR, i.e., load balancing and recursive forwarding inside the target region.

To evaluate load balancing in GEAR, SCADDS compares GEAR with GPSR, where next hop selection is completely based on geographical distance, therefore no load balancing is involved. In the traffic scenario, the team tested (1 traffic pair with source and target randomly selected, which stands for non-uniform traffic distribution), GEAR can deliver 20% to 60% more packets than GPSR before the network is partitioned. Moreover, compared to GPSR, GEAR extends network lifetime in up to 150% for small packets (180B).

Load balancing in GEAR maximizes its gain with small packets as explained. The threshold triggered beacon update will generate more overhead than GPSR. On the other hand, in order to do perimeter forwarding, GPSR carries extra fields in its header. For small packets, GEAR's extra beacon overhead and GPSR's packet overhead will cancel each other out, so that load balancing in GEAR will show its gain. However, for larger packets, the extra header fields in GPSR can be ignored, thus the beacon overhead will offset the benefits of load balancing in GEAR. However, SCADDS expects sensor network traffic to consist of mostly packets of up to a few hundred bytes due to its energy constraints. Furthermore, the beacons in the routing layer can be reduced by piggybacking them in data packets or in MAC layer control packets.

SCADDS also used ns simulations to study if recursive forwarding provides any gain over controlled flooding in the target region and how this gain changes with packet length. Unicast packet has RTS/CTS/ACK broadcast packets overhead. However, if packets are large enough, so that control overhead can be ignored, unicast recursive forwarding is still more efficient than flooding in the target region. Simulation results show that recursive forwarding delivers from 125% to 150% more packets than controlled flooding inside the target region. When MAC control overhead is ignored (large data packets), recursive forwarding exhibits major gains over controlled flooding.

3. Sensor-MAC Performance in Multi-hop Networks

S-MAC is a MAC protocol specifically designed for wire- less sensor networks developed by the SCADDS project. Building on contention-based protocols like 802.11, S-MAC strives to retain the flexibility of contention-based protocols while improving energy efficiency in multi-hop networks. S-MAC includes approaches to reduce energy consumption from all the major sources of energy waste: idle

listening, collision, overhearing and control overhead.

To measure the performance of S-MAC in multi-hop networks, the project measured the energy consumption, latency and throughput of S-MAC in a 10-hop network where nodes are placed along a line with the first node being the source and the last one being the sink.

S-MAC introduces periodic sleep on each node where neighboring nodes coordinate on their sleep schedules. The energy consumption can be greatly reduced by avoiding long-time idle listening. However, latency will be increased on each hop due to the sleep. One improvement SCADDS has made over the original S-MAC is to utilize an adaptive listening on each node. The basic idea is as follows. If a node overhears an RTS or CTS from one of its neighbors, it will wake up for a short period of time at the end of the transmission. So if the node is the next-hop node, its previous-hop neighbor will be able to pass data to it without being delayed by the normal sleep schedule.

To compare the latency in different configurations, SCADDS has made S-MAC to run in different modes. The first mode is to let each node strictly follow its sleep schedule. In the second mode, each node follows its sleep schedule but adaptively listens at the end of its neighbors' transmissions. The third mode completely disables the periodic sleep, so that there is no sleep-introduced delay on each hop.




Figure 2: Average packet latency on each hop in the lowest traffic load

Figure 2 shows the measured average latency on each packet in the lowest traffic load. Compared with the mode that each node strictly follows its sleep schedule (no adaptive listen), the mode with adaptive listen significantly reduced the latency in the multi-hop network.

SCADDS has also measured the energy consumption and throughput in this 10-hop network. The results of energy consumption basically conform to the results

the team obtained in the INFOCOM paper. Adaptive listen improves the energy efficiency in high-traffic load. It also improves the throughput in all traffic conditions compared with the mode without adaptive listen.

During the S-MAC measurement, the team also improved the robustness of node synchronization in multi-hop networks. Specifically, SCADDS added new mechanisms for nodes to discover other nodes that follow different schedules. The team is currently preparing a journal paper including all the improvement and new results of S-MAC.

4. Application Techniques

4.1 Clustering

Early SCADDS work explored adaptive clustering based techniques for an object tracking sensor network. In this application, sensors monitor location and status of various tagged objects (e.g., projectors, cameras). Queries from users about these tagged objects are efficiently resolved through the proposed schemes.

4.2 Aggregation

SCADDS examined the problem of aggregation of data as it propagates down established gradients. Aggregation can be expressed as local transformation of data to reduce overall power consumption. Hence, it might take various forms: from compressing data, to merging data temporally or spatially (i.e. data coming from different neighbors), to only forwarding deductions on the basis of received data.

Of particular interest is the approach currently being pursued to realize adaptive aggregation through the use of reinforcement. In particular, Amit is investigating an approach in which as data aggregates in a node, that node reinforces the strength of its interests accordingly and thereby pulls down additional data. Over time, aggregation points will draw additional data toward them, increasing their aggregation capability.

Preliminary results indicate that diffusion can achieve significant energy savings even with the simplest form of application-level data processing (i.e., duplicate suppression) in the network.

In-network data aggregation is essential for wireless sensor networks where resources (e.g., bandwidth, energy) are limited. In the previous approach, data is opportunistically aggregated at intermediate nodes on a low-latency tree, which may not necessarily be energy efficient. A more energy-efficient tree is a greedy tree, which can be incrementally constructed by connecting each source to the closest point of the existing tree. The current scheme (a greedy approach) constructs a greedy aggregation tree to improve path sharing.

The instantiation of directed diffusion described in the earlier work establishes low-latency paths between sources (sensor nodes that detect phenomena) and sinks (user nodes) using only localized algorithms. Paths from different sources to a sink form an aggregation tree rooted at the sink. Data from different sources is opportunistically aggregated. Whenever similar data happens to meet at a

branching node in the tree, the copies of similar data are replaced by a single message. Energy-wise, opportunistic aggregation on a low-latency tree is not optimal because data may not be aggregated (or reduced) near the sources. SCADDS proposes using a greedy incremental tree (GIT) to improve path sharing for more energy savings. The team has implemented this greedy-tree approach in ns-2 and will compare it to the prior opportunistic approach (Section 3). Due to space constraints, some of the more detailed algorithms, simulations, and analysis have been omitted. Please refer to [Intanagonwiwat02a] for more details. The preliminary results suggest that, under investigated scenarios, greedy aggregation can achieve up to 45% energy savings over opportunistic aggregation in high-density networks without adversely impacting latency or robustness.

4.3 Topology Control

With GAF, Geographic Adaptive Fidelity, [Xu01a] nodes that are redundant for communication as determined by geographical position turn off their radios in order to save energy. Nodes alternate having their radios on in order to accomplish load balancing. GAF uses location information and an idealized radio model to determine node equivalence. Location information may be provided by GPS or other location systems under development.

For the initial discussion, one assumes that there is no error in the location information. Even with location information it is not trivial to find equivalent nodes in an ad hoc network. Nodes that are equivalent for communication between one pair of nodes may not be equivalent for communication between a different pair of nodes. GAF addresses this problem by dividing the whole area where nodes are distributed into small "virtual grids".

A virtual grid is defined as follows: for two adjacent virtual grids A and B, all nodes in A can communicate with all nodes in B and vice versa. Thus, in each grid all nodes are equivalent for routing. For example, Figure 3 shows three virtual grids, A, B, and C. According to the definition of virtual grids, node 1 can reach any of nodes 2, 3, or 4, and nodes 2, 3, and 4 can all reach node 5. Therefore nodes 2, 3, and 4 are equivalent and two of them can go to sleep. The team sizes the virtual grid based on the nominal radio range R , farthest possible distance between two nodes in adjacent grids (since they must be able to communicate). If a virtual grid is a square with r units on a side, then the longest possible distance between nodes in adjacent grids is the length of the long diagonal connecting the two grids.




Figure 3: A Virtual Grid

In many settings, such as indoors or under trees where GPS does not work, location information is not available. The dependency on global location information thus limits GAF's usefulness. In addition, geographic proximity does not always lead to network connectivity. GAF must make very conservative

connectivity assumptions because it guesses at connectivity (based on a radio model) instead of directly measuring it. Being conservative requires more nodes to stay active than necessary, leading to less energy conservation. This motivates Cluster-based Energy Conservation (CEC), which, unlike GAF, does not rely on location information. Further, CEC itself directly and adaptively measures network connectivity and thus can find network redundancy more accurately so that more energy can be conserved.

CEC organizes nodes into overlapping clusters that are interconnected to each other by gateway nodes.




Figure 4: Compares GAF and CEC performance




Figure 5: Shows experimental performance of CEC in UCLA's testbed.

5. Infrastructure

5.1 Testbed Development


SCADDS experimented with several testbeds over the course of the project. Its main testbed was a network of 16 PC/104s. Several radios have been examined, including the Radiometrix in RPC and Berkley Motes with RFM radio.

Radiometrix radios had the advantage of including a high-level packet controller. While helpful to getting the testbed going, lack of control over the MAC protocol proved limiting. Therefore the project adopted a Berkely mote as the "network interface card" for the testbed and developed the S-MAC protocol (described in Section xxx).


The project also used a testbed of 8 WINSng 1.0 computers running Windows CE, later replaced with 8 WINSng 2.0 computers running Linux with integrated radios.




WINSng 1.0



WINSng 2.0




PC/104



Stayton

32 bit platforms



UCB Rene Mote

8bit Platform

Figure 6: Hardware platforms supporting diffusion

5.2 Testbed – Implementation and Analysis

5.2.1 Connectivity Measurements for SCADDS Testbed

Despite of a simple and elegant formula in theory, the propagation property of RF radio signal is hardly possible to capture in real world. Fluctuation and Asymmetry in connectivity are two major challenges when designing RF wireless ad hoc network systems. The experiments on several directed diffusion routing

applications indicate that the connectivity between testbed nodes may change dramatically over long time frame. It is very necessary to study the connectivity and its fluctuation in the testbed. Such study would improve the experimentation methodology and the understanding of wireless sensor network design in general.

5.2.2 Experiment Details

The testbed consists of 14 nodes distributed over two floors of an office building. Each node is equipped with a PC104 module and a PRC radio transceivers. The MAC protocol is a (simple) version of CSMA/CA. Beside of the wireless communication channel, the testbed also have wired infrastructure to interface itself to the Internet for debugging and management.

The quality of a link from node A to B are defined as throughput $t = \frac{N_r}{N_t}$, where

N_t is the number of packets sent by A and N_r is the packets received by B. Without loss of generality, packets are sent at the rate of 20 per second, and t is computed based on a 10 second test, i.e. around 200 pkt sent. The measurement is logged to a central server via the wired infrastructure. The experiment accesses the quality of different links in a round robin fashion: Each node in turn sends out packets as the sole transmitter and all the rest of the nodes receives. In this way, one can correlate the fluctuations of different links at approximately the same time if necessary. Taking account of the bookkeeping time as well as cleanup time to make sure all nodes are started in right sequence, each round for 14 node testbed lasts around 9 minutes.

The whole experiment lasts 45 hours, which is long enough to cover one real life cycle i.e. one business day in the offices. It is quite interesting to see if the fluctuation can be correlated to this cycle. For each source and sink pair, the quality log consists of 300 points.

5.2.3 Results

Figure [7] shows the connectivity from A to B as the average of all 300 measurements. To depict the connectivity between different nodes, the team visualizes the topology with different parameters. SCADDS plots those links with connectivity greater than certain threshold, showing how well the network is connected. With even very high threshold 0.9, the network is still connected. With threshold 0.75, the overlay topology is rich enough to do multiple hop wireless communication experiments.

SCADDS is interested in those links with heavy loss but positive connectivity because in general the source and sink usually cross multiple "hops" and data can be delivered over those links with less delay but less reliable than over those good single hop links. It seems that lossy links are quite common in this network. For Example, the links with less than 0.1 are spread over most portion of the network.




Figure 7: Experimentally measured connectivity in testbed

5.3 Software Developments

Software

- Radiometrix Device Drivers

This software, developed as part of the SCADDS project, was released to the public in June of 2000; for more information see:

<<http://www.circlemud.org/~jelson/software/radiometrix>>

This package contains a Linux device driver for the RPC (Radio Packet Controller) model of radio manufactured by Radiometrix. The RPC is a fairly low-power, self-contained, short-range, plug-on radio. It has been a critical part of the testbed infrastructure for implementation and validation of directed diffusion and other algorithms.

- Emlog

This software was also developed as part of the SCADDS project and released to the public in June of 2000; for more information see:

<<http://www.circlemud.org/~jelson/software/emlog>>

Emlog is a Linux kernel module that makes it easy to access the most recent (and only the most recent) output from a process. It works just like "tail -f" on a log file, except that the storage required never grows. This is very important for the logging and debugging facilities in embedded systems where there isn't enough memory or disk space for keeping complete log files, but the most recent debugging messages are sometimes needed (e.g., after an error is observed).

- S-MAC Communication Stack on Mica Motes

SCADDS developed a radio communication stack on the Mica Motes running TinyOS, developed by University of California, Berkeley. It is an alternative of the standard communication stack in the TinyOS release. It provides new features such as flexible packet format and headers, reliable transmission of variable length packets and the full S-MAC. Current releases are at <http://www.isi.edu/ilense/software/smac/index.html>

- Directed Diffusion

SCADDS made a number of releases of directed diffusion over the course of the project, including support for Linux, WINSng 2.0, and ns-2. Current releases are at <http://www.isi.edu/ilense/software/diffusion/index.html>

5.4 SensIT Collaboration/Integration

Over the course of the project the SCADDS group participated actively in SensIT integration activities. Examples include supporting field tests at SITEX00 and SITEX01, both of which used directed diffusion. They worked with MIT-LL to help define the diffusion API and insure that the MIT-LL and ISI implementations were compatible. They worked with researchers at MIT-LL, Cornell, PARC, BAE Systems, PSU and other institutions to support technology transfer of directed diffusion and other protocols.

6. Personnel

Over the course of the SCADDS project, the following personnel were involved:

Staff at USC/ISI:

- Deborah Estrin
- Padmaparna Haldar
- John Heidemann
- Ramesh Govindan
- Fabio Silva
- Wei Ye
- Jong-Suk Ahn
- Cengiz Alaettinoglu

Graduate Students at USC/ISI:

- Nirupama Bulusu
- Vladimir Bychkovskiy
- Jerry Elson
- Deepak Ganesan
- Lewis Girod
- Chalermek Intanagonwivat
- Amit Kumar
- Satish Kumar
- Kun-chan Lan
- Fred Stann

- Ya Xu
- Jerry Zhao
- Yen Yu

Collaborators:

- Phillippe Bonnet (DIKU)
- Professor Culler (UC Berkeley)
- Julia Liu (PARC)
- Ted Faber (USC/ISI)
- Richard Muntz (UCLA)
- Professor Pister (UC Berkeley)
- Jim Reich (PARC)
- Mani Srivastava (UCLA)
- Feng Zhao (PARC)
- Richard Brooks (PSU)
- Joe Reynolds (?)

7. References

7.1 White Papers/Technical Reports

W. Ye, J. Heidemann, D. Estrin, "A flexible and reliable radio communication stack on Mica motes" September 2002

[Heidemann02a]

John Heidemann, Fabio Silva, Yan Yu, Deborah Estrin, and Padmapartha Haldar. "Diffusion Filters as a Flexible Architecture for Event Notification in Wireless Sensor Networks" Technical Report N. ISI-TR-556, USC/Information Sciences Institute, April, 2002.

J. Ahn, J. Heidemann, "An adaptive FEC algorithm for Mobile Wireless Networks", USC/ISI Technical Report ISI-TR-555, March 2002.

"Geographical and Energy Aware Routing: A Recursive Data Dissemination Protocol for Wireless Sensor Networks". UCLA Computer Science Department Technical Report UCLA/CSD-TR-01-0023, May 2001. An improved version of this paper was submitted to WCNC' 2002.

Jong-Suk Ahn and John Heidemann, "SCADDS Data Diffusion APIs", by the SCADDS Project, first released to the SensIT mailing list October 1999.

7.2 Conference/Journal

M. Rahimi, H. Shah, G. Sukhatme, J. Heidemann, and D. Estrin. "Energy Harvesting in Mobile Sensor Networks", In Proceedings of the IEEE International Conference on Robotics and Automation, p. to appear. Taipei, Taiwan, IEEE, May 2003.

<http://www.isi.edu/~johnh/PAPERS/Rahimi03a.html>

Fred Stann and John Heidemann. "RMST: Reliable Data Transport in Sensor Networks" In Proceedings of the First International Workshop on Sensor Net Protocols and Applications, pp. 102-112. Anchorage, Alaska, USA, IEEE, April, 2003.

[Intanagonwiwat03a]

Chalermek Intanagonwiwat, Ramesh Govindan, Deborah Estrin, John Heidemann, and Fabio Silva. "Directed Diffusion for Wireless Sensor Networking" *ACM/IEEE Transactions on Networking*, V. 11 (N. 1), pp. 2-16, February, 2003.

[Ye03a]

W. Ye, J. Heidemann, and D. Estrin, "Medium Access Control with Coordinated, Adaptive Sleeping for Wireless Sensor Networks", Technical Report N. ISI-TR-567, USC/Information Sciences Institute, January 2003. expected to appear IEEE/ACM Transactions on Networking.

N. Bulusu, J. Heidemann, V. Bychkovskiy, D. Estrin, "Scalable, Ad-hoc, Deployable, RF-based Localization", In proceedings to Grace Hopper Celebration of Women in Computing, Vancouver, British Columbia, Canada, Institute for Women and Technology, October 2002.

D. Ganesan, D. Estrin, J. Heidemann, "DIMENSIONS: Why do we need a new Data Handling architecture for Sensor Networks?", In proceedings of the ACM Workshop on Hot Topics in Networks, p. to appear. Princeton, NJ, USA, ACM. October 2002.

[Intanagonwiwat02a]

Chalermek Intanagonwiwat, Deborah Estrin, Ramesh Govindan, and John Heidemann. "Impact of Network Density on Data Aggregation in Wireless Sensor Networks", In Proceedings of the 22nd International Conference on Distributed Computing Systems, pp. 457-458. Vienna, Austria, IEEE, July, 2002. See UCLA CSD TR-01-750 for an expanded version of this paper.

[Ye02a]

W. Ye, J. Heidemann, and D. Estrin, "An Energy-Efficient MAC protocol for Wireless Sensor Networks", In Proceedings of the IEEE Infocom, pp. 1567-1576, New York, NY, USA, USC/Information Sciences Institute, IEEE, June 2002.

[Zhao02a]

F. Zhao, J. Shin, and J. Reich, "Information-Driven Dynamic Sensor Collaboration for Tracking Applications", *IEEE Signal Processing Magazine*, p. to appear March 2002.

http://www.parc.xerox.com/spl/projects/cosense/pub/ieee_spm.pdf

[Zhao02b]

J. Zhao, R. Govindan, and D. Estrin, "Residual Energy Scans for Monitoring Wireless Sensor Networks", In Proceedings of the IEEE Wireless Communications and Networking Conference, pp. 356-362. Orlando, FL, USA, IEEE, March 2002.

(An extended version is available as USC CSD TR-01-745).

<http://www.isi.edu/scadds/papers/escan-wcnc02.ps.gz>

N. Bulusu, J. Heidemann, V. Bychkovskiy, D. Estrin, "Density-adaptive beacon placement algorithms for localization in ad hoc wireless networks", July 2001. Submitted to Infocom 2002.

[Heidemann01c]

J. Heidemann, F. Silva, C. Intanagonwiwat, R. Govindan, D. Estrin, D. Ganesan. "Building Efficient Wireless Sensor Networks with Low-Level Naming." In proceedings of the Symposium on Operating Systems Principles, pp. 146-159. Chateau Lake Louise, Alberta, Canada, ACM. October 2001.

[Bulusu02a]

N. Bulusu, J. Heidemann, D. Estrin, "Tradeoffs in location-support systems: The case for quality-expressive location models for applications", In proceedings of Ubicomp 2001 Workshop on Location Modeling for Ubiquitous Computing, Atlanta, GA, September 2001.

(Bulusu01c)

N. Bulusu, D. Estrin, L. Girod, and J. Heidemann, "Scalable Coordination for Wireless Sensor Networks: Self-Configuring Localization Systems", In Proceedings of the 6th IEEE International Symposium on Communication Theory and Application, St. Martin's College, Ambleside, Cumbria, UK, IEEE. July 2001.

[Xu01a]

Y. Xu, J. Heidemann, and D. Estrin, "Geography-informed Energy Conservation for Ad Hoc Routing", In Proceedings of the ACM International Conference on Mobile Computing and Networking, pp. 70-84, Rome, Italy, ACM, July 2001.

[Elson01a]

Jeremy Elson and Deborah Estrin, "Random, Ephemeral Transaction Identifiers in Dynamic Sensor Networks". In Proceedings of the International Conference on Distributed Computing Systems, pp 459-468, Phoenix, Arizona, USA, IEEE. April, 2001.

[Yu01a]

Y. Yu, R. Govindan, and D. Estrin, "Geographical and Energy Aware Routing: A Recursive Data Dissemination Protocol for Wireless Sensor Networks", Technical Report N. TR-01-0023, University of California, Los Angeles, Computer Science Department, 2001.

<http://lecs.cs.ucla.edu/~yanyu/geo-draft.ps>

[Kumar00a]

S. Kumar, C. Alaettinoglu, and D. Estrin. "SCalable Object-tracking through Unattended Techniques (SCOUT)", In Proceedings of the IEEE International Conference on Network Protocols, Osaka, Japan, IEEE. November 2000.

<http://netweb.usc.edu/estrin/papers/scout.ps.gz>

[Xu00a]

Y. Xu, J. Heidemann, D. Estrin, "Adaptive Energy-Conserving Routing for Multi-hop Ad Hoc Networks", Submitted to Mobicom 2000, August 2000, Boston, MA. USC/ISI technical paper #527, October 2000.

[Heidemann00d]

J. Heidemann, N. Bulusu, J. Elson, C. Intanagonwiwat, K. Lan, Y. Xu, W. Ye, D. Estrin, R. Govindan, "Effects of Details in Wireless Network Simulation", Presented at SCS conference on Distributed Systems modeling and Simulation Conference. pp 3-11, September 2000

[Intanagonwiwat00a]

C. Intanagonwiwat, R. Govindan, D. Estrin, "Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks", to ACM Mobicom 2000, pp 56-67

D. Estrin, R. Govindan, J. Heidemann, "Next Century Challenges: Scalable Coordination in Sensor Networks", Proceedings of ACM Mobicom 1999, Seattle, WA., pp 205-220

R. Govindan, J. Heidemann, T. Faber, D. Estrin, "Ad-Hoc Smart Environments", Proceedings of DARPA/NIST Smart Space Workshop, Atlanta, GA.